SYSTEMS, METHODS, AND MEDIA FOR HIERARCHICAL PROGRESSIVE POINT CLOUD RENDERING

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is based on, claims the benefit of, and claims priority to U.S. Provisional Application No. 62/514,706, filed Jun. 2, 2017, which is hereby incorporated herein by reference in its entirety for all purposes.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH

This invention was made with government support under HS033659 awarded by the Agency for Healthcare Research and Quality and 2016-U-CX-0017 awarded by the US Department of Justice. The government has certain rights in the invention.

BACKGROUND

Point cloud data created (e.g., using lidar techniques, which is sometimes referred to as LIDAR, LiDAR, or LADAR) that preserves the positions of objects within a scene or scenes in great detail. For example, point cloud data can be used to represent the contents of a room, a building, a crime scene, a landmark, a land formation, etc., in detail such a replica of the scene can be reproduced. Such point cloud representations may be useful for a variety of applications, such as studying details of a crime scene without blocking off the area for an extended period of time (e.g., days or weeks), for preserving sites of cultural importance (e.g., buildings, landmarks, etc.) that may otherwise be lost due to neglect, erosion, etc. As another example, point cloud data can be used to recreate the details of a patient's dwelling prior to discharge from a medical facility, such that a medical practitioner (e.g., a doctor, a physical therapist, etc.) can inspect the dwelling to determine whether any modifications should be made before the patient returns home.

While technology for creating highly detailed point cloud data that can include millions to billions of individual points is available, rendering the point cloud data can be extremely difficult in light of the potential for extremely large datasets. For example, to create a single 1920 by 1080 pixel image (e.g., an HD image) from a point cloud including just one billion points, a system must select up to about two million points from the billion available points in order to render the single image. In many virtual reality applications frame rates in the range of 70 to 120 frames per second (fps) are often required to avoid causing users discomfort. Extending the above example, to render HD quality images at 70 fps may require sifting through several dozen million points each fraction of a second (e.g., in about 14.3 milliseconds) between frames to determine which points to render, and which to ignore or discard.

One conventional technique for attempting to render point cloud data more efficiently is to organize the point cloud data in a multi-level octree that is organized based on location within the scene. The information from the octree can be used in combination with information about the field of view into the scene to relatively efficiently disregard large portions of the point cloud data that are not within the field of view. However, while this can reduce the number of points that to be considered, there are still potentially many millions of points that are within the field of view. Another conventional technique for attempting to render point cloud data more efficiently is to reduce the resolution of the image being rendered by using a single point to render multiple pixels. While this can increase frame rates at which point cloud data can be rendered, this also creates lower quality images.

Using conventional techniques to attempt to render point cloud data can result in low resolution images, low frame rates, partially transparent objects (e.g., where an object between the camera and another object is only partially rendered), blank spaces where no points are rendered, and other artifacts that can reduce the usability of the point cloud data.

Accordingly, new systems, methods, and media, for hierarchical progressive point cloud rendering are desirable.

SUMMARY

In accordance with some aspects of the disclosed subject matter, systems, methods, and media for hierarchical progressive point cloud rendering are provided.

In accordance with some embodiments of the disclosed subject matter, a method for point cloud rendering is provided, the method comprising: rendering a first image based at least in part on point cloud data, the first image corresponding to a first viewing frustum; determining that a viewing angle has changed to a second viewing frustum; requesting, from a memory, points from a first synthetic point cloud that are within the second viewing frustum, wherein the first synthetic point cloud has a resolution that is lower than a resolution of the point cloud data; requesting points in a subdivision of a second synthetic point cloud that intersects the second viewing frustum, wherein the second synthetic point cloud has a resolution that is lower than a resolution of the point cloud data and higher than a resolution of the first synthetic point cloud; reprojecting points used during rendering of the first image into a plurality of frame buffer objects of different resolutions including at least a first frame buffer object of a first resolution for points from the first synthetic point cloud, and a second frame buffer object of a second resolution for points from the second synthetic point cloud, wherein the second resolution is higher than the first resolution; determining, for each frame buffer object, whether at least one point received from the memory corresponds to the same position in the frame buffer object as a point reprojected into the frame buffer object; overwriting, in each frame buffer object, at least one reprojected point with a point corresponding to the same position received from the memory that is closer to a camera position; determining that a first position in the first frame buffer object is not associated with a reprojected point and is not associated with a point received from the memory after rendering of the first image; in response to determining that the first position in the first frame buffer object is not associated with a reprojected point and not associated with a new point, copying a point that originated in a lower resolution frame buffer object to the first position of the first frame buffer object; determining that the first frame buffer object is filled; and in response to determining that the first frame buffer object is filled, rendering a second image based on the contents of the first frame buffer object.

In some embodiments, the method further comprises: determining that a plurality of subdivisions of the second synthetic point cloud intersect the second viewing frustum; and selecting the subdivision from the plurality of subdivisions based on a distance to the subdivision.